

DIRICHLET FLOW MATCHING WITH APPLICATIONS TO DNA DESIGN

Hannes Stark*, Bowen Jing*, Chenyu Wang, Gabriele Corso
Bonnie Berger, Regina Barzilay & Tommi Jaakkola
Massachusetts Institute of Technology

ABSTRACT

Discrete diffusion or flow models could enable faster and more controllable sequence generation than autoregressive models. We show that naïve linear flow matching on the simplex is insufficient toward this goal since it suffers from discontinuities in the training target and further pathologies. To overcome this, we develop *Dirichlet flow matching* on the simplex based on mixtures of Dirichlet distributions as probability paths. In this framework, we derive a connection between the mixtures’ scores and the flow’s vector field that allows for classifier and classifier-free guidance. Further, we provide distilled Dirichlet flow matching, which enables one-step sequence generation with minimal performance hits, resulting in $O(L)$ speedups compared to autoregressive models. On complex DNA sequence generation tasks, we demonstrate superior performance compared to all baselines in distributional metrics and in achieving desired design targets for generated sequences. Finally, we show that our guidance approach improves unconditional generation and can generate DNA that satisfies design targets.

1 INTRODUCTION

Flow matching (FM) is a generative modeling framework that provides a simulation-free means of training continuous normalizing flows between noise and data distributions (Lipman et al., 2022; Liu et al., 2022; Albergo & Vanden-Eijnden, 2022) and generalizes diffusion models (Song et al., 2021). Existing formulations have yet to treat *discrete categorical data*—a notable shortcoming considering the many important applications such as biological sequence design (Avdeyev et al., 2023).

In this work, we introduce *Dirichlet flow matching* for generative modeling of discrete categorical data. We frame such modeling as a transport problem between a uniform density over the probability simplex and a finitely supported distribution over the vertices of the simplex. Naïvely, it may appear more straightforward to define a noising process that linearly interpolates between data and noise, as is the dominant approach in flow matching on \mathbb{R}^n (Liu et al., 2022; Lipman et al., 2022; Pooladian et al., 2023). However, we show that such an approach—which we call linear flow matching—suffers from pathological behavior due to the contracting support of the resulting conditional probability paths. We carefully engineer Dirichlet FM to avoid these shortcomings while admitting several advantages properties such as conditioning on a target class via *guidance* class (Dhariwal & Nichol, 2021; Ho & Salimans, 2022). Additionally, we can *distill* the sampling process into a student model and generate arbitrarily long sequences in a single forward pass.

We evaluate Dirichlet FM on three DNA sequence datasets with several complex tasks that pose diverse challenges. First, we demonstrate that Dirichlet FM better generates promoter DNA sequences with desired regulatory activity compared to baselines. Second, Dirichlet FM improves in capturing DNA data distributions with an FBD (distributional similarity) of 1.9 vs. 36.0 in Melanoma DNA and 1.0 vs. 25.2 in Fly Brain DNA. Third, we demonstrate that Dirichlet FM guidance can improve unconditional sequence generation and generate cell-type specific enhancer DNA sequences that improve upon the (experimentally validated) sequences of Taskiran et al. (2023). Lastly, distilled Dirichlet FM generates sequences in a single step, resulting in orders-of-magnitude speedups with minimal performance degradation.

*equal contribution: hstark@mit.edu, bjing@mit.edu




Figure 1: **Overview of Dirichlet flow matching.** *Left:* starting from uniform noise on the probability simplex, we define conditional probability paths that approach a point mass at the vertex via a one-parameter family of Dirichlet distributions. We view a sequence of tokens as a sequence of simplices for which the probability path corresponds to noising tokens via *superposition* with all other possible tokens (during inference, the simplices depend on each other through a joint denoiser). *Right:* Comparison of the marginal probability paths and vector fields of Dirichlet and linear FM. The vector fields of Dirichlet FM are smooth in time and space, unlike linear FM.

2 METHOD

Flow Matching on the Simplex. Let S_K be the probability simplex in K -dimensional space: $S_K = \{\mathbf{x} = (x_1, \dots, x_K)^T \in \mathbb{R}^K \mid \mathbf{1}^T \mathbf{x} = 1, \mathbf{x} \geq 0\}$. Given a K -class categorical distribution with probabilities p_i , $\sum_{i=1}^K p_i = 1$, we relax this distribution into continuous space by converting it to a mixture of point masses at the vertices of S_k (with \mathbf{e}_i as the i th one-hot vector): $p_{\text{data}}(\mathbf{x}) = \sum_{i=1}^K p_i \delta(\mathbf{x} - \mathbf{e}_i)$. We then define the noisy prior to be the uniform density on the simplex, or a *Dirichlet distribution* with parameter vector α given by the all ones vector: $q_0(\mathbf{x}) = \text{Dir}(\mathbf{x}; \alpha = (1, \dots, 1)^T) = \Gamma(K)$. Our objective is then to learn a vector field, using some choice of conditional probability path (discussed later), to transport q_0 to p_{data} . Typically, the neural network directly parameterizes the vector field and is trained via the L_2 -like conditional flow-matching loss $\mathcal{L}(\theta) = \mathbb{E}[\|u_t(\mathbf{x} \mid \mathbf{x}_1) - \hat{v}(\mathbf{x}, t; \theta)\|^2]$ where the expectation is taken over $t \sim \mathcal{U}[0, 1]$, $\mathbf{x}_1 \sim p_{\text{data}}$, $\mathbf{x} \sim p_t(\cdot \mid \mathbf{x}_1)$. However, we instead train a *denoising classifier* via a cross-entropy loss $\mathcal{L}(\theta) = \mathbb{E}[\log \hat{p}(\mathbf{x}_1 \mid \mathbf{x}; \theta)]$. At inference time, we then parameterize the vector field via

$$\hat{v}(\mathbf{x}, t; \theta) = \sum_{i=1}^K u_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i) \hat{p}(\mathbf{x}_1 = \mathbf{e}_i \mid \mathbf{x}; \theta) \quad (1)$$

It can be shown (Appendix A) that the two losses have the same minimizer, and thus, the cross-entropy is a valid flow-matching objective. The advantages of this approach are twofold: (1) it ensures that the learned vector field is restricted to the tangent plane of the simplex (i.e., the components sum to zero), and (2) the conditional vector field does not need to be evaluated at training time.

For simplicity, our discussion focuses on modeling categorical data on the simplex. However, in practice, we are interested in *sequences* of variables relaxed onto the *multi-simplex* S_K^N . At inference time, the simplices depend on each other through a learned denoiser that outputs token-wise logits conditioned on all noisy inputs.

2.1 DESIGNING SIMPLEX FLOW MATCHING

See Section C.1 for background on flow matching. As it outlines, there are two options to define a conditional probability path $p_t(\mathbf{x} \mid \mathbf{x}_1)$ and corresponding vector field $u_t(\mathbf{x} \mid \mathbf{x}_1)$ to train a flow model. **Interpolant perspective:** Define an interpolant $\psi_t(\mathbf{x}_0 \mid \mathbf{x}_1)$, which provides the density $p_t(\mathbf{x} \mid \mathbf{x}_1)$ implicitly but allows one to easily sample from it, and obtain the conditional vector field trivially by taking the derivative $u_t = \partial_t \psi_t$. **Probability path perspective:** Define $p_t(\mathbf{x} \mid \mathbf{x}_1)$ explicitly and solve for $u_t(\mathbf{x} \mid \mathbf{x}_1)$ that satisfies the transport equation which can be non-trivial.

Following extant works on flow matching, the most natural way to proceed for the simplex would be to follow the *interpolant perspective* and use the linear flow map employed in Lipman et al. (2022);

Pooladian et al. (2023) $\psi_t(\mathbf{x}_0 | \mathbf{x}_1) = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$ and its vector field $u_t(\mathbf{x} | \mathbf{x}_1) = \frac{\mathbf{x}_1 - \mathbf{x}}{1 - t} = \mathbf{x}_1 - \mathbf{x}_0$. Since S_K is a Euclidean space, these operations remain well-defined, and the interpolant transports all points on the simplex to \mathbf{x}_1 at $t = 1$ via straight paths. However, a pathological property emerges when such conditional probability paths are marginalized over $p_{\text{data}}(\mathbf{x}_1)$ over the course of flow matching training:

Proposition 1. *Suppose that a flow matching model is trained with the linear flow map. Then, for all $k = 2, \dots, K$ and $\mathbf{x} \sim p_t(\mathbf{x})$, the converged model posterior $p(\mathbf{x}_1 | \mathbf{x}) \propto p_t(\mathbf{x} | \mathbf{x}_1)p_{\text{data}}(\mathbf{x}_1)$ has support over at most $k - 1$ vertices for times $t > 1/k$.*

Conceptually, this means that as the model transports samples on the simplex at $t = 0$ to the vertices of the simplex at $t = 1$, it must eliminate or rule out a possible destination vertex at each of the times $\frac{1}{K}, \frac{1}{K-1}, \dots, \frac{1}{2}$, if not earlier. As K becomes large, an increasingly large fraction of the model capacity must be allocated to a smaller and smaller fraction of the total time and trajectory length—indeed, for all K , the posterior for times $t > 1/2$ reduces to the arg max operator. Further, the marginal field is increasingly discontinuous. We posit—and empirically verify in Section 3—that these factors significantly hurt the performance of linear FM. Upon examination of the marginal probability paths (Figure 2), it becomes clear that this pathological behavior is due to the shrinking support of the conditional paths that arise from linear FM. To address these issues, we instead turn to the *probability path perspective* and develop Dirichlet FM in the following.



Figure 2: **Pathological behavior of linear flow matching.** Each color represents a conditional probability path evolving over time toward its target vertex. At $t = 1/4$, $t = 1/3$, and $t = 1/2$, the region of overlap between 4, 3, and 2 conditional probability paths disappears, respectively, corresponding to a shrinking set of possible values of $\mathbf{x}_1 | \mathbf{x}$ for any \mathbf{x} .

2.2 DIRICHLET FLOW MATCHING

Probability path $p_t(\mathbf{x} | \mathbf{x}_1)$. Following the *probability path perspective*, we first define a conditional probability path with $t \in [0, \infty)$ as:

$$p_t(\mathbf{x} | \mathbf{x}_1 = \mathbf{e}_i) = \text{Dir}(\mathbf{x}; \boldsymbol{\alpha} = \mathbf{1} + t \cdot \mathbf{e}_i) \quad (2)$$

When $t = 0$, this is equal to the uniform noise distribution. As $t \rightarrow \infty$, the i th entry of the $\boldsymbol{\alpha}$ parameter vector increases while the others remain constant, concentrating the density towards a point mass on the i th vertex, corresponding to the $t = 1$ boundary condition in standard flow matching.¹ Hence, this family of Dirichlet distributions provides a conditional probability path with the required boundary conditions while retaining support over the entire simplex, as desired.

Vector field $u_t(\mathbf{x} | \mathbf{x}_1)$. Since we have chosen a conditional probability path directly rather than implicitly via an interpolant, it is more difficult to obtain the corresponding conditional vector field $u_t(\mathbf{x} | \mathbf{x}_1)$. Indeed, there is an infinite number of such fields that generate the desired evolution of p_t . Motivated by the basic form of the linear FM, we generalize it via the following *ansatz*:

$$u_t(\mathbf{x} | \mathbf{x}_1 = \mathbf{e}_i) = C(x_i, t)(\mathbf{e}_i - \mathbf{x}) \quad (3)$$

That is, the flow still points directly towards the target vertex \mathbf{e}_i , but is rescaled by a x_i -dependent factor. The conditional vector field of linear FM satisfies this form with $C(x_i, t) = 1/(1 - t)$ dependent only on t ; we introduce the additional x_i -dependence to control the contraction of probability mass towards \mathbf{e}_i . In Appendix A.2, we derive the $C(x_i, t)$, which gives rise to Dirichlet probability paths to be:

$$C(x_i, t) = -\tilde{I}_{x_i}(t + 1, K - 1) \frac{\mathcal{B}(t + 1, K - 1)}{(1 - x_i)^{K-1} x_i^t} \quad (4)$$

¹We continue to call the data sample \mathbf{x}_1 , and in practice, integrate to some large fixed time (typically $t = 8$) and take the arg max of the final model posterior.




Figure 3: **Vector field magnitudes** of the conditional flow field $u_t(\mathbf{x} | \mathbf{x}_1 = \mathbf{e}_i)$ as a function of x_i (the i th element of \mathbf{x}) plotted for varying values of t . In Dirichlet FM, the field vanishes at both $x_i = 1$ (i.e., the target vertex) and $x_i = 0$ (the opposite face).

where $\tilde{I}_x(a, b) = \frac{\partial}{\partial a} I_x(a, b)$ is a derivative of the *regularized incomplete beta function* $I_x(a, b)$. Figure 6 compares the magnitude of the resulting vector field with that of linear FM as a function of distance from the target vertex. As anticipated, the field vanishes both at the target vertex *and* on the $(K - 2)$ -dimensional face directly opposite it instead of monotonically scaling with distance from the target vertex as in linear FM. This second condition means the probability density is never fully drawn away from the face and resolves the pathological behavior of linear FM. The resulting probability paths and vector fields are visualized on the simplex in Figure 1; they are continuous and smooth, unlike in linear FM.

2.3 GUIDANCE

A key attribute of iterative generative models is the ability to continuously and gradually bias the generative process towards a class label with user-specified strength, a technique known as *guidance* (Dhariwal & Nichol, 2021; Ho & Salimans, 2022). Initially proposed in the context of diffusion models, where the generative process follows the *score* $\hat{s}(\mathbf{x}, t; \theta) \approx \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ of the noisy data distribution, guidance is implemented by taking a linear combination of the unconditional and conditional score models $\hat{s}_{\text{CFG}}(\mathbf{x}, t, y; \theta) = \gamma \hat{s}(\mathbf{x}, t, y; \theta) + (1 - \gamma) \hat{s}(\mathbf{x}, t, \emptyset; \theta)$ with $\gamma > 0$ and running the generative process with this adjusted score.

Relationship between flow and score. For the Dirichlet marginal probability path, the score can be obtained from the model posterior via the denoising score-matching identity (Song & Ermon, 2019) $\hat{s}(\mathbf{x}, t; \theta) = \sum_{i=1}^K s_t(\mathbf{x} | \mathbf{x}_1 = \mathbf{e}_i) \hat{p}(\mathbf{x}_1 = \mathbf{e}_i | \mathbf{x}; \theta)$. We can differentiate Equation 2 to obtain a matrix equation $\hat{s} = \mathbf{D} \hat{p}$, where $\mathbf{D}_{ij} = \delta_{ij} \frac{t}{x_i}$. Here, \mathbf{D} is a $K \times K$ diagonal matrix dependent on \mathbf{x}, t and $\hat{s}, \hat{p} \in \mathbb{R}^n$. (Technically, \hat{s} contains both on-simplex and off-simplex components, the latter of which is ignored.) Meanwhile, the computation of the marginal flow (Equation 1) can also be written as a very similar matrix equation $\hat{v} = \mathbf{U} \hat{p}$ where the entries of \mathbf{U} are given by Equation 15. Combining these, we obtain $\hat{v} = \mathbf{U} \mathbf{D}^{-1} \hat{s}$ where \mathbf{D} is invertible since it is diagonal with nonnegative entries. Thus, a *linear relationship* exists between the marginal flow and the score arising from the same model posterior.

Classifier-free guidance. Suppose we have class-conditional and unconditional flow models $\hat{v}(\mathbf{x}, t, y; \theta)$ and $\hat{v}(\mathbf{x}, t, \emptyset; \theta)$. Since a linear combination of scores results in a linear combination of flows, we similarly implement guidance by integrating $\hat{v}_{\text{CFG}}(\mathbf{x}, t, y; \theta) = \gamma \hat{v}(\mathbf{x}, t, y; \theta) + (1 - \gamma) \hat{v}(\mathbf{x}, t, \emptyset; \theta)$. **Classifier guidance.** In cases where a conditional flow model is unavailable, we use the gradient of a noisy classifier to obtain a conditional score from an unconditional score $\hat{s}(\mathbf{x}, t, y; \theta) = \hat{s}(\mathbf{x}, t, \emptyset; \theta) + \nabla_{\mathbf{x}} \log \hat{p}(y | \mathbf{x}, t; \theta)$. The conditional scores can then be converted into a model posterior \hat{p} and then a marginal flow \hat{v} via $\hat{v} = \mathbf{U} \mathbf{D}^{-1} \hat{s}$ (see Appendix A).

2.4 DISTILLATION

The aim of *distillation* (Salimans & Ho, 2022; Song et al., 2023; Yin et al., 2023) is to reduce the inference time of the iterative generative process by reducing the number of steps while retaining sample quality. However, for discrete diffusion models (see Section C.2) or autoregressive language models, no distillation techniques exist. For Dirichlet FM, inference is a deterministic ODE integration defining a map between the prior and target distribution. Hence, we can distill the teacher model (using 100 steps in our experiments) into a student model representing the map. For this, we sample the teacher to obtain pairs of noise and training targets to supervise the student model.

3 EXPERIMENTS

Promoter DNA sequence design. We next assess the ability of Dirichlet FM to design DNA promoter sequences conditioned on a desired *promoter profile*. The experimental setup and evaluation closely follow that of DDSM (Avdeyev et al., 2023), for which data details are in Appendix B.2 and implementation details in Appendix B.1. Following Avdeyev et al. (2023), we evaluate generated sequences with the mean squared error (MSE) between their predicted regulatory activity and that of the original sequence corresponding to the input profile. The regulatory activity is determined by the promoter-related predictions of SEI (Chen et al., 2022). We compare Dirichlet FM with linear FM, discrete diffusion methods, and a language model that autoregressively generates the base pairs.

Dirichlet FM improves upon Linear FM, which are the only two methods that outperform the language model baseline (Table 1). The second best method in this comparison is the distilled version of Dirichlet FM, which retains almost the same performance. This means that our distilled Dirichlet FM outperforms all other methods in a single step, which is a $100\times$ speedup compared to the diffusion models and a $1,024\times$ speedup compared to the language model.

3.1 ENHANCER DNA DESIGN

We now assess the performance of Dirichlet FM on DNA enhancer sequences and design evaluations that quantify both unconditional and conditional sample quality. Implementation and architecture details are in Appendix B.1. We evaluate on two enhancer sequence datasets from fly brain cells (Janssens et al., 2022) and from human melanoma cells (Atak et al., 2021) with labels for their cell

Table 1: **Transcription profile conditioned promoter DNA sequence generation.** The MSE is between the predicted regulatory activity of the designed sequence and the ground truth sequence. NFE is the number of function evaluations required for sampling. Numbers with * are from Avdeyev et al. (2023).

METHOD	MSE	NFE
BIT DIFFUSION*	.0395	100
D3PM-UNIFORM*	.0375	100
DDSM*	.0334	100
LANGUAGE MODEL	.0333	1024
LINEAR FM	.0281	100
DIRICHLET FM	.0269	100
DIRICHLET FM DISTILLED	.0278	1

Metric. To score the similarity between a data distribution and a generative model’s distribution, we employ a metric similar to the Fréchet inception distance (FID) that is commonly used to evaluate image generative models (Heusel et al., 2017). We follow this established principle and call our metric Fréchet Biological distance (FBD). Hence, we train a classifier model to predict cell types and use its hidden representations as embeddings of generated samples and data distribution samples. Then, the FBD is calculated as the Wasserstein distance between Gaussians fit to embeddings from the two distributions (10k each).

Q1: How well can Dirichlet FM capture the sequence distribution?

We compare with an autoregressive language model (the best baseline in the promoter design experiments in Section 3) and with Linear FM. To evaluate, we calculate the FBD between the models’ generated sequences and the unconditional data distribution. Dirichlet FM outperforms the language model by a large margin on both datasets and linear FM for human melanoma cell enhancer generation (Table 2). Moreover, distillation minimally impacts FBD while speeding up inference by *3 orders* of magnitude compared to the language model and 2 to Dirichlet FM (distilled Dirichlet FM only requires 1 step). Such speedups, compared to autoregressive models, make Dirichlet FM a promising direction for other applications with high sequence lengths where inference times are important.

Q2: Can guided Dirichlet FM produce class-specific sequences and improve upon the state-of-the-art? We condition Dirichlet FM on different target cell-type classes via guidance (Section 2.3). To quantify how well the generated sequences match the target class distribution, we use the FBD between the generated distribution and the data distribution *conditioned* on the target class. Additionally, we train a separate cell-type classifier and evaluate the probability it assigns to the target class for a generated sequence.

Sequences of classifier-free guided Dirichlet FM for the cell-type perineurial glia (PNG) have better FBD (Figure 4) than the sequences of Taskiran et al. (2023),

Table 2: **Evaluation of unconditional enhancer generation.** Each method generates 10k sequences, and we compare their empirical distributions with the data distribution using our Fréchet Biological distance (FBD) analogous to FID for image generative models. NFE refers to number of function evaluations. The RANDOM SEQUENCE baseline shows the FBD for the same number and length of sequences with uniform randomly chosen nucleotides. DIRICHLET FM DIST. refers to our one-step distilled model and DIRICHLET FM CFG to classifier-free guidance towards all classes with their empirical frequencies.

METHOD	HUMAN	FLY	NFE
	FBD	FBD	
RANDOM SEQUENCE	622.8	876	–
LANGUAGE MODEL	36.0	25.2	500
LINEAR FM	19.6	15.0	100
DIRICHLET FM	5.3	15.2	100
DIRICHLET FM DIST.	6.1	15.8	1
DIRICHLET FM CFG	1.9	1.0	200

of which several were *experimentally validated* as functioning enhancers (we show this comparison only for the PNG class since their sequences are available for it). For the other classes, guidance is similarly effective; by increasing the guidance factor γ , the classifier probability for generated sequences to belong to the target class can be improved close to 100%, and the FBD improves significantly until reaching a minimum (after the minimum the diversity decreases and the FBD worsens). The improvements with classifier guidance (Appendix Figure 11) are still significant but smaller.

Q3: Can guidance improve unconditional generation? We generate unconditional sequences with classifier-free guided Dirichlet FM by first sampling a class (based on its empirical frequency) and then guiding toward that class. This significantly improves sample quality compared to unguided Dirichlet FM (Figure 9) and baselines (Table 2). Thus, guided Dirichlet FM via the connection we derive between flow and score reproduces the success in image diffusion models of enhancing sample quality via conditioning (Rombach et al., 2022; Saharia et al., 2022). Furthermore, like guidance for images (Xu et al., 2023), the FBD first improves under increased guidance (Figure 9) until reaching a minimum, after which increased guidance deteriorates sample diversity and, therefore, FBD.

4 CONCLUSION

We presented Dirichlet flow matching for modeling discrete data via a generative process on the simplex. This solves the pathological behavior of linear flow matching on the simplex, which we identified. Compared with autoregressive methods or diffusion with discrete noise, Dirichlet FM enables distillation and conditional generation via guidance, which we derived via a connection between the marginal flow and the scores of a mixture of Dirichlets.

Experimental results on important regulatory DNA sequence design tasks across 3 datasets demonstrate Dirichlet FM’s effectiveness and utility for hard generative modeling tasks over long sequences. The results confirm Dirichlet FM’s superiority to linear FM and multiple discrete diffusion approaches. Distilling Dirichlet FM only marginally impacts performance while enabling one-step generation, leading to multiple orders of magnitude speedups compared to autoregressive generative models for long sequences. Lastly, we demonstrated effective class conditional generation via guided Dirichlet FM to design cell-type specific enhancers - an important task for gene therapies. Hence, Dirichlet FM is a flexible framework (guidance, distillation) with strong performance for biological sequence generation and a promising direction for further discrete data applications.




Figure 4: **Classifier-free guidance for cell type conditional enhancer design.** We generate enhancers that are only active in cell class via classifier-free guidance with varying γ . Shown are 4 classes of the Fly Brain cell data. The left y-axis FBD is computed between the generated sequences and the data distribution conditioned on the target class. For the first class, "PNG", functional sequences of Taskiran et al. (2023) are available, and we show their FBD. The right y-axis PROB. refers to the target class probability of a classifier for the generated sequences in percent.




Figure 5: **Classifier-free guidance for unconditional enhancer generation.** We generate unconditional sequences with classifier-free guidance by sampling the target class based on its empirical frequency. We show the FBD between the generated data and the full data distribution for varying levels of guidance γ .

REFERENCES

- Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2022.
- Zeynep Kalender Atak, Ibrahim Ihsan Taskiran, Jonas Demeulemeester, Christopher Flerin, David Mauduit, Liesbeth Minnoye, Gert Hulselmans, Valerie Christiaens, Ghanem-Elias Ghanem, Jasper Wouters, and Stein Aerts. Interpretation of allele-specific chromatin accessibility using cell state-aware deep learning. *Genome Res.*, 2021.
- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 2021.
- Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, and Jian Zhou. Dirichlet diffusion score model for biological sequence generation. *arXiv preprint arXiv:2305.10699*, 2023.
- Jason D. Buenrostro, Paul G. Giresi, Lisa C. Zaba, Howard Y. Chang, and William J. Greenleaf. Transposition of native chromatin for fast and sensitive epigenomic profiling of open chromatin, dna-binding proteins and nucleosome position. *Nature Methods*, 2013.
- Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.
- Kathleen M. Chen, Aaron K. Wong, Olga G. Troyanskaya, and Jian Zhou. A sequence-based global map of regulatory activity for deciphering human genetics. *Nature Genetics*, 2022.
- Ricky TQ Chen and Yaron Lipman. Riemannian flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2023.
- Ting Chen, Ruixiang ZHANG, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. In *The Eleventh International Conference on Learning Representations*, 2023.
- Bernardo P. de Almeida, Christoph Schaub, Michaela Pagani, Stefano Secchia, Eileen E. M. Furlong, and Alexander Stark. Targeted design of synthetic enhancers for selected tissues in the drosophila embryo. *Nature*, 2023.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- Ian Dunham, Anshul Kundaje, Shelley F. Aldred, Patrick J. Collins, Carrie A. Davis, Francis Doyle, Charles B. Epstein, Seth Fretze, Jennifer Harrow, Rajinder Kaul, Jainab Khatun, Bryan R. Lajoie, Stephen G. Landt, Bum-Kyu Lee, Florencia Pauli, Kate R. Rosenbloom, Peter Sabo, Alexias Safi, Amartya Sanyal, Noam Shores, Jeremy M. Simon, Lingyun Song, Nathan D. Trinklein, Robert C. Altshuler, Ewan Birney, James B. Brown, Chao Cheng, Sarah Djebali, Xianjun Dong, Jason Ernst, Terrence S. Furey, Mark Gerstein, Belinda Giardine, Melissa Greven, Ross C. Hardison, Robert S. Harris, Javier Herrero, Michael M. Hoffman, Sowmya Iyer, Manolis Kellis, Pouya Kheradpour, Timo Lassmann, Qunhua Li, Xinying Lin, Georgi K. Marinov, Angelika Merkel, Ali Mortazavi, Stephen C. J. Parker, Timothy E. Reddy, Joel Rozowsky, Felix Schlesinger, Robert E. Thurman, Jie Wang, Lucas D. Ward, Troy W. Whitfield, Steven P. Wilder, Weisheng Wu, Hualin S. Xi, Kevin Y. Yip, Jiali Zhuang, Bradley E. Bernstein, Eric D. Green, Chris Gunter, Michael Snyder, Michael J. Pazin, Rebecca F. Lowdon, Laura A. L. Dillon, Leslie B. Adams, Caroline J. Kelly, Julia Zhang, Judith R. Wexler, Peter J. Good, Elise A. Feingold, Gregory E. Crawford, Job Dekker, Laura Elnitski, Peggy J. Farnham, Morgan C. Giddings, Thomas R. Gingeras, Roderic Guigó, Timothy J. Hubbard, W. James Kent, Jason D. Lieb, Elliott H. Margulies, Richard M. Myers, John A. Stamatoyannopoulos, Scott A. Tenenbaum, Zhiping Weng, Kevin P. White, Barbara Wold, Yanbao Yu, John Wrobel, Brian A. Risk, Harsha P. Gunawardena, Heather C. Kuiper, Christopher W. Maier,

Ling Xie, Xian Chen, Tarjei S. Mikkelsen, Shawn Gillespie, Alon Goren, Oren Ram, Xiaolan Zhang, Li Wang, Robbyn Issner, Michael J. Coyne, Timothy Durham, Manching Ku, Thanh Truong, Matthew L. Eaton, Alex Dobin, Andrea Tanzer, Julien Lagarde, Wei Lin, Chenghai Xue, Brian A. Williams, Chris Zaleski, Maik Röder, Felix Kokocinski, Rehab F. Abdelhamid, Tyler Alioto, Igor Antoshechkin, Michael T. Baer, Philippe Batut, Ian Bell, Kimberly Bell, Sudipto Chakraborty, Jacqueline Chrast, Joao Curado, Thomas Derrien, Jorg Drenkow, Erica Dumais, Jackie Dumais, Radha Duttagupta, Megan Fastuca, Kata Fejes-Toth, Pedro Ferreira, Sylvain Foissac, Melissa J. Fullwood, Hui Gao, David Gonzalez, Assaf Gordon, Cédric Howald, Sonali Jha, Rory Johnson, Philipp Kapranov, Brandon King, Colin Kingswood, Guoliang Li, Oscar J. Luo, Eddie Park, Jonathan B. Preall, Kimberly Presaud, Paolo Ribeca, Daniel Robyr, Xiaolan Ruan, Michael Sammeth, Kuljeet Singh Sandhu, Lorain Schaeffer, Lei-Hoon See, Atif Shahab, Jorgen Skancke, Ana Maria Suzuki, Hazuki Takahashi, Hagen Tilgner, Diane Trout, Nathalie Walters, Huaian Wang, Yoshihide Hayashizaki, Alexandre Reymond, Stylianos E. Antonarakis, Gregory J. Hannon, Yijun Ruan, Piero Carninci, Cricket A. Sloan, Katrina Learned, Venkat S. Malladi, Matthew C. Wong, Galt P. Barber, Melissa S. Cline, Timothy R. Dreszer, Steven G. Heitner, Donna Karolchik, Vanessa M. Kirkup, Laurence R. Meyer, Jeffrey C. Long, Morgan Maddren, Brian J. Raney, Linda L. Grasdeder, Paul G. Giresi, Anna Battenhouse, Nathan C. Sheffield, Kimberly A. Showers, Darin London, Akshay A. Bhinge, Christopher Shestak, Matthew R. Schaner, Seul Ki Kim, Zhuzhu Z. Zhang, Piotr A. Mieczkowski, Joanna O. Mieczkowska, Zheng Liu, Ryan M. McDaniel, Yunyun Ni, Naim U. Rashid, Min Jae Kim, Sheera Adar, Zhancheng Zhang, Tianyuan Wang, Deborah Winter, Damian Keefe, Vishwanath R. Iyer, Meizhen Zheng, Ping Wang, Jason Gertz, Jost Vielmetter, E. Partridge, Katherine E. Varley, and Clarke Gasper. An integrated encyclopedia of dna elements in the human genome. *Nature*, 2012.

Griffin Floto, Thorsteinn Jonsson, Mihai Nica, Scott Sanner, and Eric Zhengyu Zhu. Diffusion on the probability simplex. *arXiv preprint arXiv:2309.02530*, 2023.

Alistair R. R. Forrest, Hideya Kawaji, Michael Rehli, J. Kenneth Baillie, Michiel J. L. de Hoon, Vanja Haberland, Timo Lassmann, Ivan V. Kulakovskiy, Marina Lizio, Masayoshi Itoh, Robin Andersson, Christopher J. Mungall, Terrence F. Meehan, Sebastian Schmeier, Nicolas Bertin, Mette Jørgensen, Emmanuel Dimont, Erik Arner, Christian Schmidl, Ulf Schaefer, Yulia A. Medvedeva, Charles Plessy, Morana Vitezic, Jessica Severin, Colin A. Semple, Yuri Ishizu, Robert S. Young, Margherita Francescato, Intikhab Alam, Davide Albanese, Gabriel M. Altschuler, Takahiro Arakawa, John A. C. Archer, Peter Arner, Magda Babina, Sarah Rennie, Piotr J. Balwierz, Anthony G. Beckhouse, Swati Pradhan-Bhatt, Judith A. Blake, Antje Blumenthal, Beatrice Bodega, Alessandro Bonetti, James Briggs, Frank Brombacher, A. Maxwell Burroughs, Andrea Califano, Carlo V. Cannistraci, Daniel Carbajo, Yun Chen, Marco Chierici, Yari Ciani, Hans C. Clevers, Emiliano Dalla, Carrie A. Davis, Michael Detmar, Alexander D. Diehl, Taeko Dohi, Finn Drabløs, Albert S. B. Edge, Matthias Edinger, Karl Ekwall, Mitsuhiro Endoh, Hideki Enomoto, Michela Fagiolini, Lynsey Fairbairn, Hai Fang, Mary C. Farach-Carson, Geoffrey J. Faulkner, Alexander V. Favorov, Malcolm E. Fisher, Martin C. Frith, Rie Fujita, Shiro Fukuda, Cesare Furlanello, Masaaki Furuno, Jun-ichi Furusawa, Teunis B. Geijtenbeek, Andrew P. Gibson, Thomas Gingeras, Daniel Goldowitz, Julian Gough, Sven Guhl, Reto Guler, Stefano Gustincich, Thomas J. Ha, Masahide Hamaguchi, Mitsuko Hara, Matthias Harbers, Jayson Harshbarger, Akira Hasegawa, Yuki Hasegawa, Takehiro Hashimoto, Meenhard Herlyn, Kelly J. Hitchens, Shannan J. Ho Sui, Oliver M. Hofmann, Ilka Hoof, Fumi Hori, Lukasz Huminiacki, Kei Iida, Tomokatsu Ikawa, Boris R. Jankovic, Hui Jia, Anagha Joshi, Giuseppe Jurman, Bogumil Kaczkowski, Chieko Kai, Kaoru Kaida, Ai Kaiho, Kazuhiro Kajiyama, Mutsumi Kanamori-Katayama, Artem S. Kasianov, Takeya Kasukawa, Shintaro Katayama, Sachi Kato, Shuji Kawaguchi, Hiroshi Kawamoto, Yuki I. Kawamura, Tsugumi Kawashima, Judith S. Kempfle, Tony J. Kenna, Juha Kere, Levon M. Khachigian, Toshio Kitamura, S. Peter Klinken, Alan J. Knox, Miki Kojima, Soichi Kojima, Naoto Kondo, Haruhiko Koseki, Shigeo Koyasu, Sarah Krampitz, Atsutaka Kubosaki, Andrew T. Kwon, Jeroen F. J. Laros, Weonju Lee, Andreas Lennartsson, Kang Li, Berit Lilje, Leonard Lipovich, Alan Mackay-sim, Ri-ichiroh Manabe, Jessica C. Mar, Benoit Marchand, Anthony Mathelier, Niklas Mejhert, Alison Meynert, Yosuke Mizuno, David A. de Lima Morais, Hiromasa Morikawa, Mitsuru Morimoto, Kazuyo Moro, Efthymios Motakis, Hozumi Motohashi, Christine L. Mummery, Mitsuyoshi Murata, Sayaka Nagao-Sato, Yutaka Nakachi, Fumio Nakahara, Toshiyuki Nakamura, Yukio Nakamura, Kenichi Nakazato, Erik van Nimwegen, Noriko Ninomiya, Hiromi Nishiyori, Shohei Noma, Tadasuke Nozaki, Soichi Ogishima, Naganari Ohkura, Hiroko Ohmiya, Hiroshi Ohno, Mitsuhiro Ohshima, Mariko Okada-Hatakeyama, Yasushi Okazaki, Valerio Orlando, Dmitry A. Ovchinnikov, Arnab Pain, Robert

- Passier, Margaret Patrikakis, Helena Persson, Silvano Piazza, James G. D. Prendergast, Owen J. L. Rackham, Jordan A. Ramilowski, Mamoon Rashid, Timothy Ravasi, Patrizia Rizzu, Marco Roncador, Sugata Roy, Morten B. Rye, Eri Saijyo, Antti Sajantila, Akiko Saka, Shimon Sakaguchi, Mizuho Sakai, Hiroki Sato, Hironori Satoh, Suzana Savvi, Alka Saxena, Claudio Schneider, Erik A. Schultes, Gundula G. Schulze-Tanzil, Anita Schwegmann, Thierry Sengstag, Guojun Sheng, Hisashi Shimoji, Yishai Shimoni, Jay W. Shin, Christophe Simon, Daisuke Sugiyama, Takaaki Sugiyama, Masanori Suzuki, Naoko Suzuki, Rolf K. Swoboda, Peter A. C. 't Hoen, Michihira Tagami, Naoko Takahashi, Jun Takai, Hiroshi Tanaka, Hideki Tatsukawa, Zuotian Tatum, Mark Thompson, Hiroo Toyoda, Tetsuro Toyoda, Eivind Valen, Marc van de Wetering, Linda M. van den Berg, Roberto Verardo, Dipti Vijayan, Ilya E. Vorontsov, Wyeth W. Wasserman, Shoko Watanabe, Christine A. Wells, Louise N. Winteringham, Ernst Wolvetang, Emily J. Wood, Yoko Yamaguchi, Masayuki Yamamoto, Misako Yoneda, Yohei Yonekura, Shigehiro Yoshida, Susan E. Zabierowski, Peter G. Zhang, Xiaobei Zhao, Silvia Zucchelli, Kim M. Summers, Harukazu Suzuki, Carsten O. Daub, Jun Kawai, Peter Heutink, Winston Hide, Tom C. Freeman, Boris Lenhard, Vladimir B. Bajic, Martin S. Taylor, Vsevolod J. Makeev, Albin Sandelin, David A. Hume, Piero Carninci, Yoshihide Hayashizaki, The FANTOM Consortium, the RIKEN PMI, and CLST (DGT). A promoter-level mammalian expression atlas. *Nature*, 2014.
- Nathan C. Frey, Daniel Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, Andreas Loukas, Vladimir Gligorijevic, and Saeed Saremi. Protein discovery with discrete walk-jump sampling. In *The Twelfth International Conference on Learning Representations*, 2024.
- Vanja Haberle and Alexander Stark. Eukaryotic core promoters and the functional basis of transcription initiation. *Nat. Rev. Mol. Cell Biol.*, 2018.
- Xiaochuang Han, Sachin Kumar, and Yulia Tsvetkov. Ssd-lm: Semi-autoregressive simplex-based diffusion language model for text generation and modular control. *arXiv preprint arXiv:2210.17432*, 2022.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 2017.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Chung-Chau Hon, Jordan A. Ramilowski, Jayson Harshbarger, Nicolas Bertin, Owen J. L. Rackham, Julian Gough, Elena Denisenko, Sebastian Schmeier, Thomas M. Poulsen, Jessica Severin, Marina Lizio, Hideya Kawaji, Takeya Kasukawa, Masayoshi Itoh, A. Maxwell Burroughs, Shohei Noma, Sarah Djebali, Tanvir Alam, Yulia A. Medvedeva, Alison C. Testa, Leonard Lipovich, Chi-Wai Yip, Imad Abugessaisa, Mickaël Mendez, Akira Hasegawa, Dave Tang, Timo Lassmann, Peter Heutink, Magda Babina, Christine A. Wells, Soichi Kojima, Yukio Nakamura, Harukazu Suzuki, Carsten O. Daub, Michiel J. L. de Hoon, Erik Arner, Yoshihide Hayashizaki, Piero Carninci, and Alistair R. R. Forrest. An atlas of human long non-coding rnas with accurate 5 ends. *Nature*, 2017.
- Ilia Igashov, Arne Schneuing, Marwin Segler, Michael Bronstein, and Bruno Correia. Retrobridge: Modeling retrosynthesis with markov bridges. In *The Twelfth International Conference on Learning Representations*, 2024.
- Jasper Janssens, Sara Aibar, Ibrahim Ihsan Taskiran, Joy N. Ismail, Alicia Estacio Gomez, Gabriel Aughey, Katina I. Spanier, Florian V. De Rop, Carmen Bravo González-Blas, Marc Dionne, Krista Grimes, Xiao Jiang Quan, Dafni Papisokrati, Gert Hulselmans, Samira Makhzami, Maxime De Waegeneer, Valerie Christiaens, Tony Southall, and Stein Aerts. Decoding gene regulation in the fly brain. *Nature*, 2022.
- Zehui Li, Yuhao Ni, Tim Huygelen, Akashaditya Das, Guoxuan Xia, Guy-Bart Stan, and Yiren Zhao. Latent diffusion model for DNA sequence generation. In *NeurIPS 2023 AI for Science Workshop*, 2023.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Yunhai Luo, Benjamin C Hitz, Idan Gabdank, Jason A Hilton, Meenakshi S Kagda, Bonita Lam, Zachary Myers, Paul Sud, Jennifer Jou, Khine Lin, Ulugbek K Baymuradov, Keenan Graham, Casey Litton, Stuart R Miyasato, J Seth Strattan, Otto Jolanki, Jin-Wook Lee, Forrest Y Tanaka, Philip Adenekan, Emma O’Neill, and J Michael Cherry. New developments on the encyclopedia of dna elements encode data portal. *Nucleic Acids Res.*, 2020.
- Anil Panigrahi and Bert W. O’Malley. Mechanisms of enhancer action: the known and the unknown. *Genome Biology*, 2021.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- Pierre H Richemond, Sander Dieleman, and Arnaud Doucet. Categorical sdes with simplex diffusion. *arXiv preprint arXiv:2210.14784*, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022.
- Toshiyuki Shiraki, Shinji Kondo, Shintaro Katayama, Kazunori Waki, Takeya Kasukawa, Hideya Kawaji, Rimantas Kodzius, Akira Watahiki, Mari Nakamura, Takahiro Arakawa, Shiro Fukuda, Daisuke Sasaki, Anna Podhajaska, Matthias Harbers, Jun Kawai, Piero Carninci, and Yoshihide Hayashizaki. Cap analysis gene expression for high-throughput analysis of transcriptional starting point and identification of promoter usage. *Proc. Natl. Acad. Sci. U. S. A.*, 2003.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *Proceedings of the 40th International Conference on Machine Learning*, pp. 32211–32252, 2023.
- Ibrahim I. Taskiran, Katina I. Spanier, Hannah Dickmanken, Niklas Kempynck, Alexandra Panıkova, Eren Can Eksi, Gert Hulselmans, Joy N. Ismail, Koen Theunis, Roel Vandepoel, Valerie Christiaens, David Mauduit, and Stein Aerts. Cell-type-directed design of synthetic enhancers. *Nature*, 2023.
- Alexander Tong, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023.
- Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. In *The Eleventh International Conference on Learning Representations*, 2023.
- Weiran Wang and Miguel A Carreira-Perpinan. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*, 2013.

Robert G. Whalen. *Promoters, Enhancers, and Inducible Elements for Gene Therapy*, pp. 60–79. Birkhäuser Boston, 1994.

Yilun Xu, Mingyang Deng, Xiang Cheng, Yonglong Tian, Ziming Liu, and Tommi S. Jaakkola. Restart sampling for improving generative processes. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T. Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. *arxiv*, 2023.

A METHOD DETAILS




Figure 6: **Vector field magnitudes** of the conditional flow field $u_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i)$ as a function of x_i (the i th element of \mathbf{x}) plotted for varying values of t . In Dirichlet FM, the field vanishes at both $x_i = 1$ (i.e., the target vertex) and $x_i = 0$ (the opposite face).

Guidance. A key attribute of iterative generative models is the ability to continuously and gradually bias the generative process towards a class label with user-specified strength, a technique known as *guidance* (Dhariwal & Nichol, 2021; Ho & Salimans, 2022). Initially proposed in the context of diffusion models, where the generative process follows the *score* $\hat{s}(\mathbf{x}, t; \theta) \approx \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ of the noisy data distribution, guidance is implemented by taking a linear combination of the unconditional and conditional score models $\hat{s}_{\text{CFG}}(\mathbf{x}, t, y; \theta) = \gamma \hat{s}(\mathbf{x}, t, y; \theta) + (1 - \gamma) \hat{s}(\mathbf{x}, t, \emptyset; \theta)$ with $\gamma > 0$ and running the generative process with this adjusted score.

Distillation. The aim of *distillation* (Salimans & Ho, 2022; Song et al., 2023; Yin et al., 2023) is to reduce the inference time of the iterative generative process by reducing the number of steps while retaining sample quality. However, for discrete diffusion models (see Section C.2) or autoregressive language models, no distillation techniques exist. For Dirichlet FM, inference is a deterministic ODE integration defining a map between the prior and target distribution. Hence, we can distill the teacher model (using 100 steps in our experiments) into a student model representing the map. For this, we sample the teacher to obtain pairs of noise and training targets to supervise the student model.

Flow Matching with Cross-Entropy Loss For all t and $\mathbf{x} \in S_K$, at convergence our denoising classifier satisfies

$$\hat{p}(\mathbf{x}_1 \mid \mathbf{x}) = \frac{p_t(\mathbf{x} \mid \mathbf{x}_1) p_{\text{data}}(\mathbf{x}_1)}{p_t(\mathbf{x})} \tag{5}$$

Thus, if we parameterize the vector field via Equation 1, then we are assured that

$$\hat{v}(\mathbf{x}, t; \theta) = \sum_{i=1}^K u_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i) \frac{p_t(\mathbf{x} \mid \mathbf{x}_1) p_{\text{data}}(\mathbf{x}_1)}{p_t(\mathbf{x})} = v(\mathbf{x}, t; \theta) \tag{6}$$

Proof of Proposition 1

Proposition 1. *Suppose that a flow matching model is trained with the linear flow map. Then, for all $k = 2, \dots, K$ and $\mathbf{x} \sim p_t(\mathbf{x})$, the converged model posterior $p(\mathbf{x}_1 \mid \mathbf{x}) \propto p_t(\mathbf{x} \mid \mathbf{x}_1) p_{\text{data}}(\mathbf{x}_1)$ has support over at most $k - 1$ vertices for times $t > 1/k$.*

Proof. In linear flow matching, the explicit form of the conditional probability path is given by

$$p_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i) = \begin{cases} [(1-t)^{K-1}\Gamma(K)]^{-1} & x_i \geq t \\ 0 & x_i < t \end{cases} \quad (7)$$

Suppose for sake of contradiction that $t > 1/k$ but $p_t(\mathbf{x} \mid \mathbf{x}_1)p_{\text{data}}(\mathbf{x}_1) \neq 0$ for $k' \geq k$ values of \mathbf{x}_1 . Without loss of generality, suppose that $\mathbf{e}_1, \dots, \mathbf{e}_k$ are k of those values. Then, by Equation 7, we have $x_i \geq t$ for $i = 1, \dots, k$. Then

$$\mathbf{1}^T \mathbf{x} = \sum_{i=1}^K x_i \geq \sum_{i=1}^k x_i \geq kt > k \cdot \frac{1}{k} = 1 \quad (8)$$

This contradicts the fact that \mathbf{x} must lie on the simplex. \square

A.1 CLASSIFIER GUIDANCE

For classifier guidance, a direct application of \hat{v} via $\hat{v} = \mathbf{U}\mathbf{D}^{-1}\hat{s}$ is not possible because the classifier gradients do not have the appropriate off-simplex components to ensure a valid model posterior (i.e., $\hat{p} \in S_K$) when operated on by \mathbf{D}^{-1} . Instead, we modify $\hat{v} = \mathbf{U}\mathbf{D}^{-1}\hat{s}$ via $\tilde{\mathbf{D}} = (\mathbf{I} - \frac{1}{K}\mathbf{1}\mathbf{1}^T)\mathbf{D}$ corresponding to projecting the score onto the tangent plane of the simplex (i.e., $\mathbf{1}^T\hat{s} = 0$). As now $\tilde{\mathbf{D}}$ is no longer invertible, we obtain the conditional model posterior from the score by solving $\hat{s} = \tilde{\mathbf{D}}\hat{p}$ with the additional constraint $\mathbf{1}^T\hat{p} = 1$ and—since the classifier in practice may result in negative probabilities—project to the simplex via the algorithm of Wang & Carreira-Perpinán (2013).

A.2 DIRICHLET CONDITIONAL VECTOR FIELD

As preliminaries, we recall the definition of the multivariate beta function:

$$\mathcal{B}(\alpha_1, \dots, \alpha_K) = \frac{\prod_{i=1}^K \Gamma(\alpha_i)}{\Gamma\left(\sum_{i=1}^K \alpha_i\right)} \quad (9)$$

and that of the incomplete (two-argument) beta function:

$$\mathcal{B}(x; a, b) = \int_0^x t^{a-1}(1-t)^{b-1} dt \quad (10)$$

with the identity $\mathcal{B}(a, b) = \mathcal{B}(1; a, b)$.

We wish to construct a conditional flow $u_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i)$ which generates the evolution of the conditional probability path

$$p_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i) = \text{Dir}(\mathbf{x}; 1, \dots, \alpha_i = 1+t, \dots, 1) = \frac{\Gamma(t+K)}{\Gamma(t+1)} x_i^t \quad (11)$$

We choose the following *ansatz* for the functional form of u_t :

$$u_t(\mathbf{x} \mid \mathbf{x}_1 = \mathbf{e}_i) = C(x_i, t)(\mathbf{e}_i - \mathbf{x}) \quad (12)$$

i.e., (1) the flow points towards the target vertex \mathbf{e}_i and (2) the magnitude is scaled by a constant dependent only on x_i and t . Then consider the $(K-2)$ -dimensional hyperplane A_b of constant $x_i = b$ which cuts through the simplex, separating it into two regions S_1, S_2 (Figure 7). We make the following key observation:

The probability flux crossing the plane A_b is equal to the the rate of change of the total probability of region S_1 .

Thus, we solve for the constant $C(x_i, t)$ in the *ansatz* by deriving these two quantities and setting them equal to each other.

Q1: What is the probability mass of S_1 and its rate of change?




Figure 7: Conceptual derivation of the conditional vector field

The probability mass of S_1 can be obtained by integrating the density over each hyperplane A_c (defined by $x_i = c$) for $c < b$ and then integrating over all such hyperplanes. Since the density is a constant proportional to $x_i^t = c^t$ over each hyperplane, we obtain

$$p_t(S_1) \propto \int_0^b c^t \cdot \text{Vol}(A_c) dc \quad (13)$$

where $\text{Vol}(A_c)$ refers to the volume of the intersection between the hyperplane and the simplex. Since this region is defined by $x_i = c$, the remaining entries of \mathbf{x} must add up to $1 - c$. Thus, A_c can be viewed as a nonstandard probability simplex over $K - 1$ variables. The volume of this $(K - 2)$ -dimensional space is proportional to $(1 - c)^{K-2}$, giving

$$\rho(S_1) \propto \int_0^b c^t (1 - c)^{K-2} dc = \mathcal{B}(b; t + 1, K - 1) \quad (14)$$

Normalizing, we obtain

$$p_t(S_1) = \frac{\int_0^b c^t (1 - c)^{K-2} dc}{\int_0^1 c^t (1 - c)^{K-2} dc} = \frac{\mathcal{B}(b; t + 1, K - 1)}{\mathcal{B}(t + 1, K - 1)} = I_b(t + 1, K - 1) \quad (15)$$

where I is the so-called *regularized incomplete beta function* and is well-known as the CDF of the Beta distribution. Its derivative with respect to the first parameter is not available in closed form, but we write it as

$$\tilde{I}_x(a, b) = \frac{\partial}{\partial a} I_x(a, b) = \frac{\partial}{\partial a} \frac{\mathcal{B}(x; a, b)}{\mathcal{B}(a, b)} \quad (16)$$

and thus obtain the rate of change of $p_t(S_1)$ as $\tilde{I}_b(t + 1, K - 1)$.

Q2: What is the probability flux across the hyperplane A_b ? The probability flux across A_b (into S_2) is given by

$$J = \iint_{A_b} u_i p_t \cdot \hat{\mathbf{n}} dA = \iint_{A_b} p_t \cdot C(b, t) (\mathbf{e}_i - \mathbf{x}) \cdot \frac{\mathbf{n}}{\|\mathbf{n}\|} dA \quad (17)$$

The normal vector \mathbf{n} points from the center of the face opposite \mathbf{e}_i , specified by $x_i = 0$ and $x_j = 1/(K - 1)$, $j \neq i$, towards \mathbf{e}_i . Thus, the probability flux density is given by the dot product of

$$\begin{aligned} \mathbf{e}_i - \mathbf{x} &= (-x_1, \dots, x_i = 1 - b, \dots, -x_K) \\ \mathbf{n} &= \left(-\frac{1}{K - 1}, \dots, x_i = 1, \dots, -\frac{1}{K - 1} \right) \end{aligned} \quad (18)$$

Since $\sum_{j \neq i} x_j = 1 - b$, this dot product is equal to $(1 - b)K/(K - 1)$. We also see that $\|\mathbf{n}\| = \sqrt{K/(K - 1)}$. Importantly, the flux density is constant on the hyperplane A , so the total flux is given by a simple product:

$$J = C(b, t) \cdot p_t \cdot (1 - b) \sqrt{\frac{K}{K - 1}} \text{Vol}(A_b) \quad (19)$$

$$= C(b, t) \cdot \left[\frac{1}{\sqrt{K}} \frac{\Gamma(t+K)}{\Gamma(t+1)} b^t \right] (1-b) \sqrt{\frac{K}{K-1}} \frac{\sqrt{K-1}}{\Gamma(K-1)} (1-b)^{K-2} \quad (20)$$

$$= C(b, t) \cdot \frac{(1-b)^{K-1} b^t}{\mathcal{B}(t+1, K-1)} \quad (21)$$

where (for this step only) we note that the volume of a simplex over K variables *viewed as a subset* of \mathbb{R}^K is $\sqrt{K}/\Gamma(K)$, giving the Dirichlet PDF an additional factor of $1/\sqrt{K}$. Now setting $J = -\partial_t p_t(S_1)$ from above, we obtain

$$C(b, t) = -\tilde{I}_b(t+1, K-1) \frac{\mathcal{B}(t+1, K-1)}{(1-b)^{K-1} b^t} \quad (22)$$

Checking the transport equation. We check that $\nabla \cdot (p_t u_t) = -\partial p_t / \partial t$:

$$\nabla \cdot (p_t u_t) = -\nabla \cdot \left[\frac{\Gamma(t+K)}{\Gamma(t+1)} x_i^t \cdot \tilde{I}_{x_i}(t+1, K-1) \frac{\mathcal{B}(t+1, K-1)}{(1-x_i)^{K-1} x_i^t} (\mathbf{e}_i - \mathbf{x}) \right] \quad (23)$$

$$= -\Gamma(K-1) \nabla \cdot \left[\frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} (\mathbf{e}_i - \mathbf{x}) \right] \quad (24)$$

$$= -\Gamma(K-1) \left[\frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} \nabla \cdot (\mathbf{e}_i - \mathbf{x}) + (\mathbf{e}_i - \mathbf{x}) \cdot \nabla \frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} \right] \quad (25)$$

The divergence of \mathbf{x} on the $(K-1)$ -dimensional space S_K is $K-1$. Also, the gradient of a function dependently only on x_i has nonzero component only in the x_i direction.

$$= -\Gamma(K-1) \left[\frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} (1-K) + (1-x_i) \frac{\partial}{\partial x_i} \frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} \right] \quad (26)$$

$$= -\Gamma(K-1) \left[\frac{\tilde{I}_{x_i}(t+1, K-1)}{(1-x_i)^{K-1}} (1-K) + (1-x_i) \left(\tilde{I}_{x_i} \frac{K-1}{(1-x_i)^K} + \frac{1}{(1-x_i)^{K-1}} \frac{\partial \tilde{I}_{x_i}}{\partial x_i} \right) \right] \quad (27)$$

The first and second terms now cancel.

$$= -\Gamma(K-1) \left[\frac{1}{(1-x_i)^{K-2}} \frac{\partial}{\partial x_i} \tilde{I}_{x_i}(t+1, K-1) \right] \quad (28)$$

Substituting Equation 16 and interchanging the order of derivatives,

$$= -\Gamma(K-1) \left[\frac{1}{(1-x_i)^{K-2}} \frac{\partial}{\partial t} \frac{1}{\mathcal{B}(t+1, K-1)} \frac{\partial}{\partial x_i} \mathcal{B}(x_i; t+1, K-1) \right] \quad (29)$$

$$= -\frac{1}{(1-x_i)^{K-2}} \frac{\partial}{\partial t} \frac{\Gamma(K-1)}{\mathcal{B}(t+1, K-1)} x_i^t (1-x_i)^{K-2} \quad (30)$$

$$= -\frac{\partial}{\partial t} \frac{\Gamma(t+K)}{\Gamma(t+1)} x_i^t = -\frac{\partial p_t}{\partial t} \quad (31)$$

B EXPERIMENTAL DETAILS

B.1 TRAINING AND INFERENCE

Toy experiments. We train all models in Figure 10 for 450,000 steps with a batch size of 512 to ensure that they have all converged and then evaluate the KL of the final step.

Promoter Design. We follow the setup of Avdeyev et al. (2023) and train for 200 epochs with a learning rate of 5×10^{-4} and early stopping on the MSE on the validation set. We communicated

with Avdeyev et al. (2023) to ensure that we have the same training and inference setup as them, and we build on their codebase to evaluate the generated sequences with the Sei regulatory activity prediction model (Chen et al., 2022). Thus, we use 100 inference steps for our Dirichlet FM instead of the 400 that they use in their code since they state that they used 100 integration steps for the results in the paper, which is also stated in the paper. Under this setup, we obtain the performance numbers for Linear FM and Dirichlet FM. We also ran the autoregressive language model in this setup (except for the number of inference steps which does not apply). Meanwhile, the numbers we report for Bit Diffusion, D3PM, and DDSM are taken from the DDSM paper (Avdeyev et al., 2023).

Enhancer Design. For both evaluations, on the human melanoma cell and the fly brain cell dataset, we train for 800 epochs (convergence of validation curves is reached after approximately 300 for both datasets). We use FBD for early stopping. For inference, we use 100 integration steps.

Classifier-free Guidance. For classifier-free guidance, we train with a conditioning ratio (the fraction of times we train with a class label as input instead of the no-class token as input) of 0.7. During inference, the inference Algorithm is changed in that two probability vectors \hat{p} are predicted, once with class conditioning and once without class conditioning, which we sum together according to the classifier free guidance equation. Then, we project the resulting probabilities onto the simplex since negative values can arise for guidance factors $\gamma > 1$. For this purpose, we use the algorithm by Wang & Carreira-Perpinán (2013).

Classifier Guidance. The noisy classifier that we train for classifier guidance has the same architecture as our generative model, except that we sum the final representations and feed them into a 2-layer feed-forward network that serves as classification head. For training, we use early stopping on the accuracy and train for approximately 800 epochs. During inference, we use automatic differentiation to obtain the classifier’s gradients with respect to the input points on the simplices. To perform classifier guidance, we then convert the flow model output probabilities to scores as described in Section 2.3, obtain the guided score by adding the unconditional and the conditional scores, and convert the obtained scores back to probabilities. These we project to the simplex (Wang & Carreira-Perpinán, 2013) from which we obtain the vector field for integration.

Classifier for FBD calculation. This classifier’s architecture is similar to that of the noisy classifier for classifier guidance. However, it does not have any time conditioning and takes token embeddings as input instead of points on the simplex. For training, we use early stopping on the accuracy and train approximately 100 epochs.

The sequence embeddings that we use to calculate FBD are given by the hidden features after the first layer of the classification head. The 4 classes that we choose for the cell type specific enhancer generation are chosen as classes with a good tradeoff between the area under the receiver operator characteristic curve and the area under the precision-recall curve.

Distillation. For distillation, we run inference with the teacher model for every training step of the student model to obtain pairs of noise and training targets. With this, we train the student model on approximately 6 billion sequences.

Computational requirements. We train on RTX A600 GPUs. Training in the enhancer generation setup for 200 epochs on sequences with length 500 takes 7 hours. Our largest model uses approximately 8GB of RAM during training.

Architecture The architecture that we use for the promoter design experiments is the same as in DDSM (Avdeyev et al., 2023). In our other experiments we replaced group norm with layer norm in their architecture. The model consists of 20 layers of 1D convolutions interleaved with time embedding layers (and class-type embedding layers for classifier-free guidance) and normalization layers. We also experimented with Transformer architectures, which led to worse performance. All models use this 20-layer architecture except for the classifier for the Fly Brain data, which has 5 layers.

B.2 DATA

Promoter Data. We use a dataset of 100,000 promoter sequences with 1,024 base pairs extracted from a database of human promoters (Hon et al., 2017). Each sequence has a CAGE signal (Shiraki et al., 2003) annotation available from the FANTOM5 promoter atlas (Forrest et al., 2014), which




Figure 8: Histograms of cell type classes for fly brain cell enhancer sequence data and human melanoma cell data.

indicates the likelihood of transcription initiation at each base pair ($\mathbf{r} \in \mathbb{R}^{1024}$). Sequences from chromosomes 8 and 9 are used as a test set, and the rest for training.

For the enhancer data of 104665 fly brain cell sequences (Janssens et al., 2022), we use the same split as Taskiran et al. (2023), resulting in an 83726/10505/10434 split for train/val/test. Meanwhile, for the human melanoma cell dataset of 88870 sequences (Atak et al., 2021), their split has 70892/8966/9012 sequences. It is noteworthy that these datasets contain ATAC-seq data (Buenrostro et al., 2013), which means that not all sequences are guaranteed to be enhancers and actually enhance transcription of a certain gene. ATAC-seq only measures the chromatin accessibility of the sequences in the cell types, which is a necessary but not sufficient requirement for a sequence to be an enhancer. In Figure 8, we show histograms for the class distributions of both datasets. For the melanoma dataset, there is little class imbalance.

C BACKGROUND

C.1 FLOW MATCHING

In flow matching (Lipman et al., 2022; Liu et al., 2022; Albergo & Vanden-Eijnden, 2022), we consider a noisy distribution $\mathbf{x}_0 \sim q_0$ and data distribution $\mathbf{x}_1 \sim p_{\text{data}}$ and regress a neural network against a vector field that transports $q(\mathbf{x}_0)$ to $p_{\text{data}}(\mathbf{x}_1)$. To do so, we define a *conditional probability path*—a time-evolving distribution $p_t(\mathbf{x} | \mathbf{x}_1)$, $t \in [0, 1]$ conditioned on \mathbf{x}_1 with boundary conditions $p_0(\mathbf{x} | \mathbf{x}_1) = q(\mathbf{x})$ and $p_1(\mathbf{x} | \mathbf{x}_1) \approx \delta(\mathbf{x} - \mathbf{x}_1)$. We additionally assume knowledge of a *conditional vector field* $u_t(\mathbf{x} | \mathbf{x}_1)$ that generates $p_t(\mathbf{x} | \mathbf{x}_1)$, i.e. satisfying the transport equation

$$\frac{\partial}{\partial t} p_t + \nabla \cdot (p_t u_t) = 0 \quad (32)$$

Then, the marginal probability path

$$p_t(\mathbf{x}) = \int p_t(\mathbf{x} | \mathbf{x}_1) p_{\text{data}}(\mathbf{x}_1) d\mathbf{x}_1 \quad (33)$$

interpolates between noise $p_0 = q_0$ and data $p_1 \approx p_{\text{data}}$ and is generated by the *marginal vector field*

$$v_t(\mathbf{x}) = \int u_t(\mathbf{x} | \mathbf{x}_1) \frac{p_t(\mathbf{x} | \mathbf{x}_1) p_{\text{data}}(\mathbf{x}_1)}{p_t(\mathbf{x})} d\mathbf{x}_1 \quad (34)$$

Thus, by learning and integrating a neural network $\hat{v}(\mathbf{x}, t; \theta) \approx v_t(\mathbf{x})$, we can generate data from noisy samples $\mathbf{x}_0 \sim q$. The core design decision is the choice of appropriate conditional probability path $p_t(\mathbf{x} | \mathbf{x}_1)$ and associated vector field $u_t(\mathbf{x} | \mathbf{x}_1)$. Although it is possible to define these directly, it is often simpler to instead define a *conditional flow map* $\psi_t(\mathbf{x}_0 | \mathbf{x}_1)$ that directly transports $\mathbf{x}_0 \sim q$ to the intermediate distribution $p_t(\mathbf{x} | \mathbf{x}_1)$. The flow map immediately provides the corresponding vector field:

$$u_t(\mathbf{x} | \mathbf{x}_1) = \frac{d}{dt} \psi_t(\mathbf{x}_0 | \mathbf{x}_1) \quad (35)$$

With this formulation, the required boundary conditions simplify to $\psi_0(\mathbf{x}_0 | \mathbf{x}_1) = \mathbf{x}$ and $\psi_1(\mathbf{x}_0 | \mathbf{x}_1) = \mathbf{x}_1$. As advocated by several works (Liu et al., 2022; Lipman et al., 2022; Pooladian et al., 2023; Tong et al., 2023), the flow map (also called *interpolant*) is often chosen to follow the simplest possible path between the two endpoints—e.g., linear in Euclidean spaces and geodesic on Riemannian manifolds (Chen & Lipman, 2023).

C.2 DISCRETE DIFFUSION MODELS

Existing discrete diffusion frameworks can be split into 4 categories. Firstly, simplex-based approaches frame discrete data as vertices of a simplex and generate it starting from a Dirichlet prior over the whole simplex (Richemond et al., 2022; Floto et al., 2023). Among those, DDSM (Avdeyev et al., 2023) is most related to our work and converges to a Dirichlet distribution via Jacobi diffusion processes and the stick-breaking transform. We note that none of these simplex-based approaches feature Dirichlet distributions as intermediate distributions of the noising process—a key aspect of our approach.

The second class of discrete diffusion models fully relaxes discrete data into continuous space without any constraints and uses, e.g., a standard Gaussian as prior (Han et al., 2022; Chen et al., 2023; Frey et al., 2024). The third paradigm, established by D3PM (Austin et al., 2021), operates on discrete samples of noise distributions constructed by injecting discrete noise into data (Campbell et al., 2022; Igashov et al., 2024; Vignac et al., 2023). Lastly, *latent* discrete diffusion models train an additional network to obtain continuous latents for which they train a conventional diffusion model (Dieleman et al., 2022; Li et al., 2023).

C.3 PROMOTER AND ENHANCER DNA

DNA is a sequence with *base pairs* as tokens (3 billion for humans) and a vocabulary of 4 *nucleotides* (A, T, C, G). Parts of DNA encode genes that are transcribed into mRNA and then translated to functional proteins. *Promoters* and *enhancers* refer to noncoding portions of DNA that regulate the expression level of these genes and play important roles in eukaryotic organisms such as humans (Dunham et al., 2012; Luo et al., 2020). More specifically, a promoter for a gene is the DNA sequence next to the gene where the transcriptional machinery binds and starts transcribing DNA to mRNA (Haberle & Stark, 2018). Meanwhile, enhancers are sequences that can be distant in the DNA sequence (millions of base pairs) but are close in 3D space (Panigrahi & O’Malley, 2021) and regulate the recruitment of this transcriptional machinery. Unlike promoters, enhancers often regulate transcription in specific cell types. Hence, while both types of DNA subsequences are important for gene therapy (Whalen, 1994), the cell type specificity of enhancers enables targeting, e.g., only cancer cells.

Designing enhancers. Recently Taskiran et al. (2023) and de Almeida et al. (2023) demonstrated cell type-specific enhancer design via an *optimization* procedure starting from an initial sequence guided by a cell-type activity classifier. However, such sequence designs may not follow the empirical distribution of enhancers, which would be captured by a generative model. Our work sets the foundation for more principled conditional sequence design by learning and drawing from conditional data distributions.

D ADDITIONAL RESULTS

D.1 GUIDANCE TO IMPROVE *unconditional* GENERATION

We generate unconditional sequences with classifier-free guided Dirichlet FM by first sampling a class (based on its empirical frequency) and then guiding toward that class. This significantly improves sample quality compared to unguided Dirichlet FM (Figure 9) and baselines (Table 2). Thus, guided Dirichlet FM via the connection we derive between flow and score reproduces the success in image diffusion models of enhancing sample quality via conditioning (Rombach et al., 2022; Saharia et al., 2022). Furthermore, like guidance for images (Xu et al., 2023), the FBD first improves under increased guidance (Figure 9) until reaching a minimum, after which increased guidance deteriorates sample diversity and, therefore, FBD.




Figure 9: **Classifier-free guidance for unconditional enhancer generation.** We generate unconditional sequences with classifier-free guidance by sampling the target class based on its empirical frequency. We show the FBD between the generated data and the full data distribution for varying levels of guidance γ .




Figure 10: **Scaling to higher simplex dimensions.** We train on simple categorical distributions with an increasing number of categories K and measure the KL divergence between the generated distributions (512k samples) and the training target distribution. Dirichlet FM scales to larger K much better than linear FM.

D.2 SIMPLEX DIMENSION TOY EXPERIMENT

We first evaluate Dirichlet FM and linear FM in a simple toy experiment where the KL divergence of the generated distribution to the target distribution can be evaluated under increasing simplex dimensions. For this, we train both methods to reproduce a categorical distribution $q(\mathbf{x}) = \text{Cat}(K, \mathbf{a})$ where the class probabilities \mathbf{a} are sampled from a uniform Dirichlet $\mathbf{a} \sim \text{Dir}(\mathbf{a}; \mathbf{1})$. To evaluate, we use the KL-divergence $\text{KL}(\tilde{p} \parallel q)$ where \tilde{p} is the empirical distribution of 512,000 model samples.

Figure 10 shows these KL-divergences for increasing sizes K of the categorical distribution. Dirichlet FM is able to overfit on the simple distribution regardless of K . Meanwhile, Linear FM is unable to overfit on simple categorical distributions as K increases, illustrating the practical implications of the pathological probability paths and discontinuous vector fields as discussed in Section 2.1.

D.3 ANALYTICAL TOY EXPERIMENT FOR CLASSIFIER GUIDANCE.

As a toy experiment to demonstrate our classifier guidance procedure, we construct a distribution conditioned on a binary random variable. The conditional distribution is a categorical distribution over 20 classes. In this setup, the time-dependent class probabilities conditioned on a noisy point on the simplex can be computed analytically. Thus, we can use them to obtain the vector field for Dirichlet FM analytically. Furthermore, the class probabilities and the gradients of their logarithm (the score) can be computed analytically. Hence, we can simulate classifier guided Dirichlet FM analytically for this toy distribution.

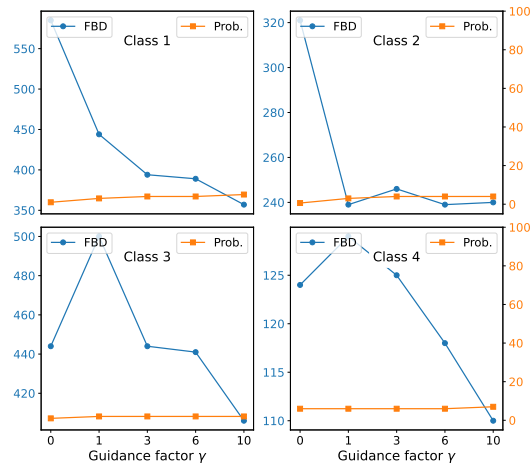


Figure 11: **Evaluation for cell type specific enhancer design with classifier guidance.** We condition Dirichlet FM to generate sequences that are only active in a class of cell type via classifier guidance with varying γ . Results are shown for 4 classes in the Fly Brain cell data. The left y-axis FBD is computed between the generated sequences and the data distribution conditioned on the target class. The right y-axis PROB. refers to the target class probability of a classifier for the generated sequences.




Figure 12: **Simulation results for classifier guidance on a conditional categorical distribution with an analytically tractable vector field.** In blue, we show a histogram (with frequencies normalized to density) of the generated classes for each class on the x-axis. In orange, we show the true probabilities of the toy distribution conditioned on the same class that we use for classifier guidance.




Figure 13: Per class performance of the classifiers used for evaluation and FBD calculation. Shown are scatter plots for each class of the fly brain data (left) and the human melanoma cell data (right) between the area under the curve of the receiver operator characteristic (x-axis) and the area under the precision-recall curve (y-axis). The scatter plots match the results of the classifiers of Atak et al. (2021) and Janssens et al. (2022) closely.

The results in Figure 12 show a close match between the empirical distribution of the generated data and the ground truth probabilities. This empirically confirms the effectiveness of our classifier guidance procedure that relies on converting probabilities to scores and converting them back to probabilities by solving a linear system of equations.